

Fast Average-Case Pattern Matching on Weighted Sequences

Carl Barton¹, Chang Liu², and Solon P. Pissis²

¹ The Blizard Institute, Barts and The London School of Medicine and Dentistry,
Queen Mary University of London, UK

c.barton@qmul.ac.uk

² Department of Informatics, King's College London, UK
{chang.2.liu,solon.pissis}@kcl.ac.uk

Abstract. A *weighted string* over an alphabet of size σ is a string in which a set of letters may occur at each position with respective occurrence probabilities. Weighted strings, also known as *position weight matrices* or *uncertain sequences*, naturally arise in many contexts. In this article, we study the problem of weighted string matching with a special focus on average-case analysis. Given a weighted *pattern* string x of length m , a *text* string y of length $n > m$, and a *cumulative weight threshold* $1/z$, defined as the minimal probability of occurrence of factors in a weighted string, we present an algorithm requiring average-case search time $o(n)$ for pattern matching for *weight ratio* $\frac{z}{m} < \min\{\frac{1}{\log z}, \frac{\log \sigma}{\log z(\log m + \log \log \sigma)}\}$. For a pattern string x of length m , a weighted text string y of length $n > m$, and a cumulative weight threshold $1/z$, we present an algorithm requiring average-case search time $o(\sigma n)$ for the same weight ratio. The importance of these results lies on the fact that these algorithms work in average-case sublinear search time in the size of the text, and in linear preprocessing time and space in the size of the pattern, for these ratios.

1 Introduction

A weighted string over some alphabet is a type of *uncertain* sequence in which a *set* of letters (instead of a single letter) may occur at each position with respective occurrence probabilities. This notion was first introduced by Iliopoulos et al. [11,12] in 2003. A great deal of research has been conducted ever since on weighted strings: for pattern matching [13,2,4]; for computing various types of regularities [14,7,18,5]; for indexing [2,10]; and for alignments [3].

An *alphabet* Σ is a finite non-empty set of size σ , whose elements are called *letters*. A *string* on an alphabet Σ is a finite, possibly empty, sequence of elements of Σ . The zero-letter sequence is called the *empty string*, and is denoted by ε . The *length* of a string x is defined as the length of the sequence associated with the string x , and is denoted by $|x|$. We denote by $x[i]$, for all $0 \leq i < |x|$, the letter at index i of x . Each index i , for all $0 \leq i < |x|$, is a position in x when $x \neq \varepsilon$. It follows that the i -th letter of x is the letter at position $i - 1$ in x .

The *concatenation* of two strings x and y is the string of the letters of x followed by the letters of y ; it is denoted by xy . A string x is a *factor* of a string

y if there exist two strings u and v , such that $y = uxv$. Consider the strings x, y, u , and v , such that $y = uxv$, if $u = \varepsilon$ then x is a *prefix* of y , if $v = \varepsilon$ then x is a *suffix* of y . Let x be a non-empty string and y be a string, we say that there exists an *occurrence* of x in y , or more simply, that x *occurs in* y , when x is a factor of y . Every occurrence of x can be characterised by a position in y ; thus we say that x occurs at the *starting position* i in y when $y[i..i + |x| - 1] = x$.

Single nucleotide polymorphisms, as well as errors introduced by wet-lab sequencing platforms during the process of DNA sequencing, can occur in some positions of a DNA sequence. In some cases, these uncertainties can be accurately modelled as a *don't care* letter. However, in other cases they can be more subtly expressed, and, at each position of the sequence, a probability of occurrence can be assigned to each letter of the nucleotide alphabet; this process gives rise to a *weighted string* or a *position weight matrix*. For instance, consider a IUPAC-encoded [1] DNA sequence, where the ambiguity letter M occurs at some position of the sequence, representing either base A or base C. This gives rise to a weighted DNA sequence, where at the corresponding position of the sequence, we can assign to each of A and C an occurrence probability of 0.5.

A weighted string x of length n on an alphabet Σ is a finite sequence of n sets. Every $x[i]$, for all $0 \leq i < n$, is a set of ordered pairs $(s_j, \pi_i(s_j))$, where $s_j \in \Sigma$ and $\pi_i(s_j)$ is the probability of having letter s_j at position i . Formally, $x[i] = \{(s_j, \pi_i(s_j)) | s_j \neq s_\ell \text{ for } j \neq \ell, \text{ and } \sum_j \pi_i(s_j) = 1\}$. A letter s_j *occurs* at position i of a weighted string x if and only if the *occurrence probability* of letter s_j at position i , $\pi_i(s_j)$, is greater than 0. A string u of length m is a *factor* of a weighted string if and only if it occurs at starting position i with *cumulative occurrence probability* $\prod_{j=0}^{m-1} \pi_{i+j}(u[j]) > 0$. Given a *cumulative weight threshold* $1/z \in (0, 1]$, we say that factor u is *valid*, or equivalently that factor u has a valid occurrence, if it occurs at starting position i and $\prod_{j=0}^{m-1} \pi_{i+j}(u[j]) \geq 1/z$. Similarly, we say that letter s_j at position i is valid if $\pi_i(s_j) \geq 1/z$. For succinctness of presentation, if $\pi_i(s_j) = 1$ the set of pairs is denoted only by the letter s_j ; otherwise it is denoted by $[(s_{j_1}, \pi_i(s_{j_1})), \dots, (s_{j_k}, \pi_i(s_{j_k}))]$.

In this article, we consider the following two problems.

WEIGHTEDPATTERNMATCHING

Input: a weighted string x of length m , a string y of length $n > m$, and a cumulative weight threshold $1/z \in (0, 1]$

Output: all positions i of y where a valid factor of length m of x occurs

WEIGHTEDTEXTMATCHING

Input: a string x of length m , a weighted string y of length $n > m$, and a cumulative weight threshold $1/z \in (0, 1]$

Output: all positions i of y where a valid factor v of y occurs and $x = v$

In [2], Amir et al. showed an $\mathcal{O}(nz^2 \log z)$ -time bound for pattern matching on weighted strings via reduction. A direct $\mathcal{O}(nz^2 \log z)$ -time algorithm for solving this problem was recently presented in [6]. Many other algorithms for solving these problems exist, but their runtime efficiency relies on the assumption of

a given *constant* z (cf. [13,10]). In this article, we are interested in designing average-case efficient algorithms for certain realistic *weight ratios* z/m .

Our Contribution. We present efficient average-case algorithms for weighted string matching. Specifically, we present two new algorithms: one to solve problem WEIGHTEDPATTERNMATCHING and another one to solve problem WEIGHTEDTEXTMATCHING. Both algorithms can achieve average-case sublinear search time in the size of the text, and work in linear preprocessing time and space in the size of the pattern. Essentially, we show that they achieve these average-case search times depending on the number of positions required, in every matching weighted factor of length m , to have a letter occurring with probability greater than $1 - 1/z$. We consider these conditions to be a quite realistic scenario in a wide range of applications, in particular, on molecular sequences [16].

2 Properties and Auxiliary Data Structures

We first start by providing a description of the key ideas on which our algorithms are based. This informal description should help the reader understand the structure and the analysis of our algorithms later on. These ideas are not all new but we find their combination in this context novel.

- We perform a colouring stage, similar to [2] and [10], on the weighted pattern string. This stage assigns a colour to each position of the weighted string based on the occurrence probabilities of the corresponding letters. We can then obtain (shown in [2] and [10]) an upper bound on the special colour of *hard* positions that any valid factor can contain.
- With this upper bound at hand, we can make an assumption to design our average-case algorithms. By ignoring these hard positions, we can search fast for candidate occurrences using only the *easy* positions (positions that are easily verified) to filter out positions of the text that could never yield a valid match. This assumption poses a first condition on the weight ratios z/m .
- For the average-case analysis of the designed algorithms, we need a randomness model on weighted strings. In the standard setting, this is trivial assuming uniform distribution of letters. For weighted strings, in order to provide a realistic scenario, we view a weighted string as an indeterminate string [9] and assume uniform distribution of subsets of the alphabet.

Given a weighted string x of length m , we perform a colouring stage on x , similar to the one before the construction of the weighted suffix tree [10], which assigns a colour to every position in x according to the following scheme:

- mark position i *black* (**b**), if *none* of the occurring letters at position i has probability of occurrence greater than $1 - 1/z$.
- mark position i *grey* (**g**), if *one* of the occurring letters at position i has probability of occurrence greater than $1 - 1/z$ and less than 1.

- mark position i *white* (\mathfrak{w}), if *one* of the occurring letters at position i has probability of occurrence 1.

Notice that if $z \geq 2$, then at every white and grey position there is only one valid letter since only one letter can have probability of occurrence greater than $1 - 1/z \geq 1/2$, whereas in a black position there maybe several valid letters. However, if $z < 2$ there are no letters with probability of occurrence at least $1 - 1/z$ in a black position since all letters have probability of occurrence of at most $1 - 1/z > 1/z$. Therefore for the rest of this article we assume $z \geq 2$. The colouring stage can be trivially performed in time $\mathcal{O}(\sigma m)$.

Lemma 1 ([2,10]). *Given a weighted string x and a cumulative weight threshold $1/z \in (0, 1]$, any valid factor of x contains at most $\ell = \lceil \log z / \log(\frac{z}{z-1}) \rceil$ black positions.*

Proof. Any letter at a black position of x has probability of occurrence less than or equal to $1 - 1/z$. For any valid factor of x , it thus holds (in the worst case) that $(1 - 1/z)^\ell \geq 1/z$. Taking the logarithm at both sides yields the lemma. \square

The second key idea of the designed algorithms comes from the following simple fact. This idea is used in many *other* pattern matching problems on strings [15].

Lemma 2. *Given a weighted string x and a cumulative weight threshold $1/z \in (0, 1]$, if $\ell < m$, then there exists a consecutive sequence of positions of length at least $\lfloor \frac{m}{\ell+1} \rfloor$ of x consisting of only white and grey positions.*

Proof. Immediate from the pigeonhole principle. \square

We can also preprocess a weighted string x of length m to compute a matrix $A[0 \dots \ell' - 1, 0 \dots \sigma - 1]$, such that for each black position i , $0 \leq i < \ell'$, and each letter $a \in \Sigma$, we have $A[i, \alpha] = 1$ if α occurs at the i th black position of x and 0 otherwise. After such a $\mathcal{O}(\sigma m)$ -time preprocessing, we can check in constant time whether a letter in a black position of x matches a letter from another string or not. With matrix A at hand, we can proceed with a fast verification step of the pre-computed candidate occurrences using Lemma 3 (see below).

Given two strings u and v in the standard setting, we say that the probability that $u[i] = v[i]$, for some position i on u and v , is given by $1/\sigma$ assuming uniform distribution of letters of the alphabet per position. This randomness model cannot be applied on weighted strings, where a subset of the alphabet occurs at every position of the string. For a given position, we rather assume a uniform distribution of all possible subsets of the alphabet, such that each letter of the subset has probability of occurrence greater than 0; i.e., such that the letter occurs. Under this assumption we can obtain the following lemma.

Lemma 3. *Given a string u and a weighted string v , the expected length of the longest valid prefix of v that is also a valid prefix of u is bounded by six.*

We start with a few definitions to reduce this problem to another one before giving a proof. An *indeterminate string* x of length m on an alphabet Σ is a finite sequence of m sets, such that $x[i] \subseteq \Sigma$, $x[i] \neq \emptyset$, for all $0 \leq i < m$. If $|x[i]| = 1$, that is, $x[i]$ represents a single letter of Σ , we say that $x[i]$ is a *solid* letter. We say that two indeterminate strings x and y *match*, denoted by $x \approx y$, if $|x| = |y|$ and for each $i = 0, \dots, |x| - 1$, we have $x[i] \cap y[i] \neq \emptyset$.

Proof (of Lemma 3). We view the weighted string v as indeterminate string v' of length $|v|$ such that $a \in v'[i]$ iff $(a, \pi(a)) \in v[i]$ and $\pi(a) > 0$. Since we completely ignore letter probabilities and thereby the validity of factors—all factors are now valid—it suffices to show that the expected number $s > 0$ of positions such that $u[0..s-1] \approx v'[0..s-1]$ and $u[s] \notin v'[s]$ is bounded by six.

We consider the comparison of u and v' from left to right. We have that $\{u[i]\} \cap v'[i] \neq \emptyset$ occurs in the following cases:

- $v'[i]$ is solid and $\{u[i]\} = v'[i]$
- $v'[i]$ is not solid and $u[i] \in v'[i]$.

Thus the total number of positive comparisons is

$$\sigma \sum_{j=1}^{\sigma} \frac{j \binom{\sigma}{j}}{\sigma} = \sum_{j=1}^{\sigma} j \binom{\sigma}{j} = \sigma 2^{\sigma-1}.$$

The total number of any case is $\sigma \sum_{j=1}^{\sigma} \binom{\sigma}{j} = \sigma(2^{\sigma} - 1)$. Therefore the probability r of $\{u[i]\} \cap v'[i] \neq \emptyset$ is

$$r = \frac{2^{\sigma-1}}{2^{\sigma} - 1} \leq 2/3, \text{ for } \sigma > 1.$$

Thus the expected number $s > 0$ of positions such that $u[0..s-1] \approx v'[0..s-1]$ and $u[s] \notin v'[s]$ can be described by the summation of infinite terms

$$s = r + 2r^2 + \dots = \sum_{k=1}^{\infty} kr^k,$$

which is bounded by $r/(1-r)^2 = 6$, for $r \leq 2/3$. This concludes the proof. \square

In [2] it was shown that $\ell = \mathcal{O}(z \log z)$. Here we refine this to an exact bound which is useful later on in the analysis of the algorithms.

Lemma 4. *Let $z \geq 2$. Then $\ell \leq z \log z$.*

Proof. By Lemma 1 we know that $\ell = \lceil \frac{\log z}{\log(\frac{z}{z-1})} \rceil$. For $z > 1$, we must show that:

$$\ell = \lceil \frac{\log z}{\log(\frac{z}{z-1})} \rceil = \lceil \frac{\log z}{\log(z) - \log(z-1)} \rceil \leq z \log z. \text{ Or equivalently that:}$$

$$\frac{\log z (z \log z - z \log(z-1) - 1)}{\log z - \log(z-1)} > 0.$$

Clearly the above is true if and only if: $z \log z - z \log(z - 1) - 1 > 0$. There is a discontinuity at $z = 1$; after this it is *always* positive and the following holds:

$$\lim_{z \rightarrow \infty} z \log z - z \log(z - 1) - 1 = 0.$$

□

3 Algorithms

We are now in a position to present Algorithm WPM to solve problem WEIGHTEDPATTERNMATCHING. In this problem, we are given a weighted string x of length m , a string y of length $n > m$, and a cumulative weight threshold $1/z \in (0, 1]$, and we need to find all positions i of y where a valid factor of length m of x occurs.

Algorithm WPM($x, m, y, n, 1/z, \Sigma$)

Perform the colouring stage on x ;
Find the number ℓ' of black positions in x ;
 $\sigma \leftarrow |\Sigma|$;
Compute $A[0 \dots \ell' - 1, 0 \dots \sigma - 1]$ of x ;
if $\ell' < m$ **then**
 Find the longest factor f in x with no black positions;
else
 return *FAIL*;
Search for f in y ;
foreach *occurrence of f in y* **do**
 Check if f is extensible to the left using A ;
 Check if f is extensible to the right using A ;
 if *the length of extension is at least m* **then**
 Verify the validity of the factor of x and report the position;

Theorem 1. *Algorithm WPM correctly solves problem WEIGHTEDPATTERNMATCHING, achieving average-case search time $o(n)$, if and only if*

$$\frac{z}{m} < \min \left\{ \frac{1}{\log z}, \frac{\log \sigma}{\log z (\log m + \log \log \sigma)} \right\}.$$

Algorithm WPM requires preprocessing time and space $\mathcal{O}(\sigma m)$.

Proof. Let $\ell < m$. By Lemma 4, the maximal value of ℓ is $z \log z$, hence when $z \log z < m$ we obtain condition $\frac{z}{m} < \frac{1}{\log z}$. In this case, by Lemma 2 and the correctness of the average-case time-optimal searching algorithm [17], all

positions i of y where a valid factor of length m of x occurs must be naïvely verified in the **for** loop of the algorithm; therefore the algorithm is correct.

The colouring stage on string x can be trivially performed in time $\mathcal{O}(\sigma m)$. The preprocessing time and space cost for array A of x is $\mathcal{O}(\sigma m)$. Assuming $\ell < m$, by Lemma 2, the minimum factor length in x with no black positions is at least $\lfloor \frac{m}{\ell+1} \rfloor$. This factor f is viewed as a standard string obtained by choosing in all grey and black positions the most probable letter. Searching f in y can be performed in average-case time $\mathcal{O}(\frac{n \log(m/\ell)}{m/\ell})$ [17]. The preprocessing time and space for searching is $\mathcal{O}(m)$. The number of expected occurrences is $n/\sigma^{\lfloor m/(\ell+1) \rfloor}$.

Let us denote the cost of verification per occurrence by $\text{VER}(m, z)$. Algorithm WPM achieves average-case search time $\mathcal{O}(\frac{n \log(m/\ell)}{m/\ell}) = o(n)$ if and only if

$$\frac{\text{VER}(m, z)n}{\sigma^{\lfloor m/(\ell+1) \rfloor}} \leq c \frac{n(\ell+1) \log_{\sigma} \frac{m}{\ell+1}}{m}$$

for some fixed constant c . That is, the total average-case verification cost is no more than the average-case searching cost. We take σ -based logarithms to obtain

$$\log_{\sigma} \text{VER}(m, z) + \log_{\sigma} m - \log_{\sigma} c - \log_{\sigma}(\ell+1) - \log_{\sigma} \log_{\sigma} \frac{m}{\ell+1} \leq m/(\ell+1).$$

By Lemma 3 and using array A of x , we know it is possible to pick $c = \text{VER}(m, z)$ to obtain a maximum value for ℓ , that is

$$\log_{\sigma} \frac{m}{(\ell+1) \log_{\sigma} \frac{m}{\ell+1}} \leq m/(\ell+1), \text{ which gives } \ell < m \left(\frac{1}{\log_{\sigma} \frac{m}{(\ell+1) \log_{\sigma} \frac{m}{\ell+1}}} \right).$$

Therefore we get the following second condition, simplified slightly for comprehension,

$$\ell < m \left(\frac{\log \sigma}{\log m - \log(\ell+1) + \log \log \sigma - \log \log m + \log \log(\ell+1)} \right).$$

By Lemma 4, $\ell \leq z \log z$, and by the previous condition on ℓ , $\ell < m$, we can further simplify the second condition to

$$\frac{z}{m} < \frac{\log \sigma}{\log z(\log m + \log \log \sigma)},$$

and this concludes the proof. \square

Corollary 1. *Let $\sigma > m \log \sigma$. Algorithm WPM achieves average-case search time $o(n)$, if and only if there exist at least 1 position in x , where a letter occurs with probability greater than $1 - 1/z$.*

Proof. For $\sigma > m \log \sigma$, we know $\log \sigma > \log m + \log \log \sigma$. By Theorem 1, Algorithm WPM achieves average-case search time $o(n)$, if and only if $\frac{z}{m} < \frac{1}{\log z}$. By Lemma 4, we know that $\ell \leq z \log z$, for all $z \geq 2$, which gives $\ell < m$. This means that there can be at most $m - 1$ black positions in x . \square

We can similarly obtain the following complementary corollary.

Corollary 2. *Let $\sigma \leq m \log \sigma$. Algorithm WPM achieves average-case search time $o(n)$, if and only if there exist at least $\lceil \frac{m(\log m + \log \log \sigma) - m \log \sigma}{\log m + \log \log \sigma} \rceil$ positions in x , where a letter occurs with probability greater than $1 - 1/z$.*

Proof. For $\sigma \leq m \log \sigma$, we know $\log \sigma \leq \log m + \log \log \sigma$. By Theorem 1, Algorithm WPM achieves average-case search time $o(n)$, if and only if $\frac{z}{m} < \frac{\log \sigma}{\log z(\log m + \log \log \sigma)}$. By Lemma 4, we know that $\ell \leq z \log z$, for all $z \geq 2$, which gives $\ell < \frac{m \log \sigma}{\log m + \log \log \sigma}$. This means that there can be at most $\lfloor \frac{m \log \sigma}{\log m + \log \log \sigma} \rfloor$ black positions in x . \square

We next present Algorithm WTM to solve problem WEIGHTEDTEXTMATCHING. In this problem, we are given a string x of length m , a weighted string y of length $n > m$, and a cumulative weight threshold $1/z \in (0, 1]$, and we need to find all positions i of y where a valid factor v of length m of y occurs and $v = x$.

For the searching stage of Algorithm WTM, we view weighted string y as the string y' according to the following scheme:

- if position i is white, then $y'[i] = \alpha$, where $(\alpha, \pi(\alpha)) \in y[i]$, $\alpha \in \Sigma$, and $\pi(\alpha) = 1$.
- if position i is grey, then $y'[i] = \alpha$, where $(\alpha, \pi(\alpha)) \in y[i]$, $\alpha \in \Sigma$, and $\pi(\alpha) > 1 - 1/z$.
- if position i is black, then $y'[i] = \lambda$, where $\lambda \notin \Sigma$.

Intuitively, while searching, we assign for every black position of y to y' a letter λ that is not in Σ . This in turn implies that writing a position on string y' requires time $\mathcal{O}(\sigma)$. In [8], it was shown that searching for a set of patterns of total length M in a text of length n requires average-case search time $\mathcal{O}(\frac{n \log m}{m})$ if the length m of the shortest pattern in the set is polynomial in M . Hence we can do searching for a set of patterns in y' in average-case time $\mathcal{O}(\frac{\sigma n \log m}{m})$. The additional factor σ is due to the cost of writing and, subsequently, reading a letter of string y' . Notice that string y' is implicit, we never actually construct it; and we never perform a colouring stage on y as these would require time $\mathcal{O}(\sigma n)$.


```

Algorithm  $WTM(x, m, y, n, 1/z, \Sigma)$ 
 $\ell \leftarrow \lceil \log z / \log(\frac{z}{z-1}) \rceil$ ;
if  $\ell < m$  then
    Partition  $x$  in  $\ell + 1$  non-overlapping fragments  $f_0, f_1, \dots, f_\ell$ ;
    Each fragment is of length at most  $\lceil \frac{m}{\ell+1} \rceil$  and at least  $\lfloor \frac{m}{\ell+1} \rfloor$ ;
else
    return FAIL;
Search for  $f_0, f_1, \dots, f_\ell$  by considering string  $y'$ ;
foreach occurrence of  $f \in \{f_0, f_1, \dots, f_\ell\}$  in  $y'$  do
    Check if  $f$  is extensible to the left naïvely;
    Check if  $f$  is extensible to the right naïvely;
    if the length of extension is at least  $m$  then
        Verify the validity of the factor of  $y$  and report the position;

```

Theorem 2. *Algorithm WTM correctly solves problem WEIGHTEDTEXTMATCHING, achieving average-case search time $o(\sigma n)$, if and only*

$$\frac{z}{m} < \min \left\{ \frac{1}{\log z}, \frac{\log \sigma}{\log z (\log m + \log \log \sigma)} \right\}.$$

Algorithm WTM requires preprocessing time and space $\mathcal{O}(m)$.

Proof. Let $\ell < m$. By Lemma 4, the maximal value of ℓ is $z \log z$, hence when $z \log z < m$ we obtain condition $\frac{z}{m} < \frac{1}{\log z}$. In this case, by Lemma 2 and the correctness of the average-case time-optimal searching algorithm [8], all positions i of y where a valid factor v of y occurs, such that $x = v$, must be naïvely verified in the **for** loop of the algorithm; therefore the algorithm is correct.

Assuming $\ell < m$, by Lemma 2, the minimum factor length with no black positions in any factor of length m of y is at least $\lfloor \frac{m}{\ell+1} \rfloor$. Searching stage of the $\ell + 1$ fragments of x (of length at most $\lceil \frac{m}{\ell+1} \rceil$ and at least $\lfloor \frac{m}{\ell+1} \rfloor$) in string y' can be performed in average-case search time $\mathcal{O}(\frac{\sigma n \log(m/\ell)}{m/\ell})$ (see [8] and the analysis above). The preprocessing time and space for searching is $\mathcal{O}(m)$. The number of expected occurrences is $n(\ell + 1)/\sigma^{\lfloor m/(\ell+1) \rfloor}$.

Let us denote the cost of verification per occurrence by $\sigma \text{VER}(m, z)$. The additional factor σ is due to the cost of reading a letter of string y naïvely. Algorithm WTM achieves average-case search time $\mathcal{O}(\frac{\sigma n \log(m/\ell)}{m/\ell}) = o(\sigma n)$ if and only if

$$\frac{(\ell + 1)\sigma \text{VER}(m, z)n}{\sigma^{\lfloor m/(\ell+1) \rfloor}} \leq c \frac{\sigma n (\ell + 1) \log \sigma \frac{m}{\ell+1}}{m}$$

for some fixed constant c . That is, the total average-case verification cost is no more than the average-case searching cost. We take σ -based logarithms to obtain

$$\log_\sigma \text{VER}(m, z) + \log_\sigma m - \log_\sigma c - \log_\sigma \log_\sigma \frac{m}{\ell+1} \leq m/(\ell+1).$$

By Lemma 3, we know it is possible to pick $c = \text{VER}(m, z)$ to obtain

$$\log_\sigma m - \log_\sigma \log_\sigma \frac{m}{\ell+1} \leq m/(\ell+1).$$

This gives a maximum value for ℓ , that is

$$\log_\sigma \frac{m}{\log_\sigma \frac{m}{\ell+1}} \leq m/(\ell+1), \text{ which gives } \ell < m \left(\frac{1}{\log_\sigma \frac{m}{\log_\sigma \frac{m}{\ell+1}}} \right).$$

Therefore we get the following second condition, simplified slightly for comprehension,

$$\ell < m \left(\frac{\log \sigma}{\log m + \log \log \sigma - \log \log m + \log \log(\ell+1)} \right).$$

By Lemma 4, $\ell \leq z \log z$, and the previous condition on ℓ , $\ell < m$, we can further simplify the second condition to

$$\frac{z}{m} < \frac{\log \sigma}{\log z(\log m + \log \log \sigma)},$$

and this concludes the proof. \square

Similar to the previous problem we obtain the following complementary corollaries.

Corollary 3. *Let $\sigma > m \log \sigma$. Algorithm WTM achieves average-case search time $o(\sigma n)$, if and only if there exist at least 1 position in any factor of length m of y , where a letter occurs with probability greater than $1 - 1/z$.*

Corollary 4. *Let $\sigma \leq m \log \sigma$. Algorithm WTM achieves average-case search time $o(\sigma n)$, if and only if there exist at least $\lceil \frac{m(\log m + \log \log \sigma) - m \log \sigma}{\log m + \log \log \sigma} \rceil$ positions in any factor of length m of y , where a letter occurs with probability greater than $1 - 1/z$.*

In Table 1, we summarise the results presented in this section and conclude with the following remarks.

Remark 1. In most typical applications we have that $\sigma \leq m \log \sigma$, hence, Corollaries 2 and 4 apply rather than Corollaries 1 and 3, respectively.

A typical example application of Corollaries 2 and 4 is provided below.

Example 1. Let $m = 32$ and $\sigma = 4$. Algorithms WPM and WTM achieve search time sublinear, on average, in the size of the text, if every matching weighted factor of length m has at most 10 black (uncertain) positions.

| Algo. | Weighted String | Weight Ratio (z/m) | Space | Preprocess. Time | Search Time |
|-------|-----------------|--|-------------------------|-------------------------|----------------|
| WPM | Pattern | $\min \left\{ \frac{1}{\log z}, \frac{\log \sigma}{\log z (\log m + \log \log \sigma)} \right\}$ | $\mathcal{O}(\sigma m)$ | $\mathcal{O}(\sigma m)$ | $o(n)$ |
| WTM | Text | $\min \left\{ \frac{1}{\log z}, \frac{\log \sigma}{\log z (\log m + \log \log \sigma)} \right\}$ | $\mathcal{O}(m)$ | $\mathcal{O}(m)$ | $o(\sigma n)$ |

Table 1: Average-case algorithms for weighted string matching

Remark 2 (Worst-Case). If Algorithm WPM returns FAIL, we can naïvely verify all positions of y in worst-case time $\mathcal{O}(\sigma m + nz^2 \log z)$ which matches the worst-case search time of the algorithm. Similarly the worst-case search time for Algorithm WTM is $\mathcal{O}(\sigma n + nz^2 \log z)$.

4 Final Remarks

Amir et al., in their seminal work [2], showed an $\mathcal{O}(nz^2 \log z)$ -time bound for pattern matching on weighted strings via reduction. A direct $\mathcal{O}(nz^2 \log z)$ -time algorithm for solving this problem was recently presented in [6]. This algorithm is efficient when z is constant or logarithmic in the length of the pattern.

In this article, we designed algorithms for weighted string matching that can achieve search time sublinear, on average, in the size of the text. We also showed exact bounds on the weighted ratio z/m in order to achieve this time complexity. Both the cost of preprocessing time and that of space requirements are linear in the size of the pattern.

As a by-product, we also showed upper bounds on when the designed algorithms achieve these average-case search times. These bounds depend on the number of positions required, in every matching weighted factor of length m , to have a letter occurring with probability greater than $1 - 1/z$ (Corollaries 1–4). We consider these conditions to be a quite realistic scenario in a wide range of applications. A particular example application is for finding IUPAC-encoded nucleotide or peptide sequences such as *cis*-elements in nucleotide sequences or small domains and motifs in protein sequences [16].

Our immediate target is twofold.

1. From a theoretical point of view, we are planning to extend our approach here to address the problem of weighted string matching when both the pattern and the text are weighted strings.
2. From a practical point of view, we are planning to implement the presented average-case algorithms and the worst-case algorithm presented in [6], and evaluate them with real and synthetic data.

Acknowledgements

We warmly thank G. Fici from the University of Palermo for useful discussions.

References

1. Nomenclature Committee of the International Union of Biochemistry (NC-IUB). Nomenclature for incompletely specified bases in nucleic acid sequences. Recommendations 1984. *European Journal of Biochemistry* 150(1), 1–5 (1985)
2. Amir, A., Chencinski, E., Iliopoulos, C.S., Kopelowitz, T., Zhang, H.: Property matching and weighted matching. *Theor. Comput. Sci.* 395(2-3), 298–310 (2008)
3. Amir, A., Gotthilf, Z., Shalom, B.R.: Weighted LCS. *Journal of Discrete Algorithms* 8(3), 273–281 (2010)
4. Amir, A., Iliopoulos, C., Kapah, O., Porat, E.: Approximate matching in weighted sequences. In: Lewenstein, M., Valiente, G. (eds.) *Combinatorial Pattern Matching, Lecture Notes in Computer Science*, vol. 4009, pp. 365–376. Springer Berlin Heidelberg (2006)
5. Barton, C., Iliopoulos, C.S., Pissis, S.P.: Optimal computation of all tandem repeats in a weighted sequence. *Algorithms for Molecular Biology* 9(21) (2014)
6. Barton, C., Pissis, S.P.: Linear-time computation of prefix table for weighted strings. In: Manea, F., Nowotka, D. (eds.) *Combinatorics on Words - 10th International Conference, WORDS 2015, Kiel, Germany, September 14-17, 2015, Proceedings. Lecture Notes in Computer Science*, vol. 9304, pp. 73–84. Springer (2015)
7. Christodoulakis, M., Iliopoulos, C.S., Mouchard, L., Perdikuri, K., Tsakalidis, A.K., Tsichlas, K.: Computation of repetitions and regularities of biologically weighted sequences. *Journal of Computational Biology* 13(6), 1214–1231 (2006)
8. Crochemore, M., Czumaj, A., Gasieniec, L., Lecroq, T., Plandowski, W., Rytter, W.: Fast practical multi-pattern matching. *Inf. Process. Lett.* 71(3-4), 107–113 (1999)
9. Holub, J., Smyth, W.F., Wang, S.: Fast pattern-matching on indeterminate strings. *J. Discrete Algorithms* 6(1), 37–50 (2008)
10. Iliopoulos, C.S., Makris, C., Panagis, Y., Perdikuri, K., Theodoridis, E., Tsakalidis, A.K.: The weighted suffix tree: An efficient data structure for handling molecular weighted sequences and its applications. *Fundam. Inform.* 71(2 and 3), 259–277 (2006)
11. Iliopoulos, C.S., Mouchard, L., Perdikuri, K., Tsakalidis, A.K.: Computing the repetitions in a weighted sequence. In: Simánek, M. (ed.) *Proceedings of the Prague Stringology Conference 2003, Prague, Czech Republic, September 22-24, 2003*. pp. 91–98. Department of Computer Science and Engineering, Faculty of Electrical Engineering, Czech Technical University (2003)
12. Iliopoulos, C.S., Mouchard, L., Perdikuri, K., Tsakalidis, A.K.: Computing the repetitions in a biological weighted sequence. *Journal of Automata, Languages and Combinatorics* 10(5/6), 687–696 (2005)
13. Iliopoulos, C.S., Pedikuri, K., Tsichlas, K.: The pattern matching problem in biological weighted sequences. In: Feraggina, P., Grossi, R. (eds.) *Proc. of Third International Conference on Fun with Algorithms (FUN 2004)*. pp. 106–117 (2004)

14. Iliopoulos, C.S., Perdikuri, K., Theodoridis, E., Tsakalidis, A.K., Tsihlias, K.: Motif extraction from weighted sequences. In: Apostolico, A., Melucci, M. (eds.) String Processing and Information Retrieval, 11th International Conference, SPIRE 2004, Padova, Italy, October 5-8, 2004, Proceedings. Lecture Notes in Computer Science, vol. 3246, pp. 286–297. Springer (2004)
15. Navarro, G.: A guided tour to approximate string matching. *ACM Comput. Surv.* 33(1), 31–88
16. Yan, T., Yoo, D., Berardini, T.Z., Mueller, L.A., Weems, D.C., Weng, S., Cherry, J.M., Rhee, S.Y.: PatMatch: a program for finding patterns in peptide and nucleotide sequences. *Nucleic Acids Research* 33(Suppl 2), W262–W266 (2005)
17. Yao, A.C.C.: The complexity of pattern matching for a random string. *SIAM J. Comput.* 8(3), 368–387 (1979)
18. Zhang, H., Guo, Q., Iliopoulos, C.S.: Locating tandem repeats in weighted sequences in proteins. *BMC Bioinformatics* 14(S-8), S2 (2013)